Cosmic Ray Statistics using LabVIEW
AAPT Summer 3013, Workshop W36

David B. Pengra
Department of Physics
University of Washington
Seattle, Washington

Apparatus

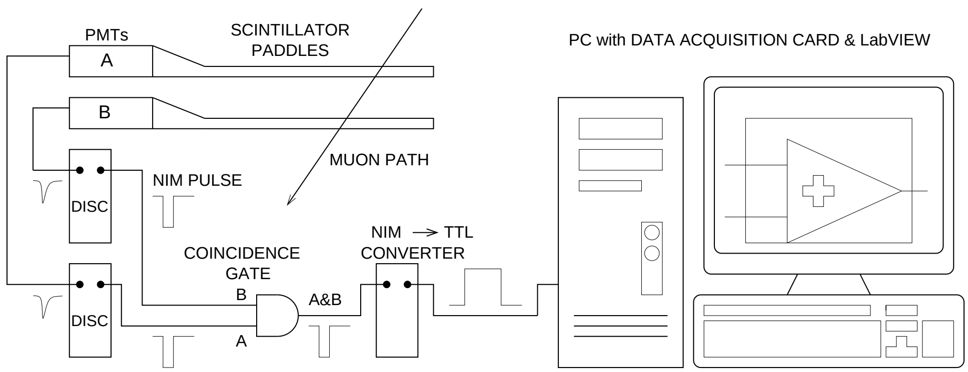

Figure 1: Schematic of the muon telescope apparatus. Although there are only two scintillator paddles shown in the figure, the actual setup uses between three and four paddles. The AND gate is a 4-fold NIM logic unit. Data are collected with the LabVIEW Interval Counter software.

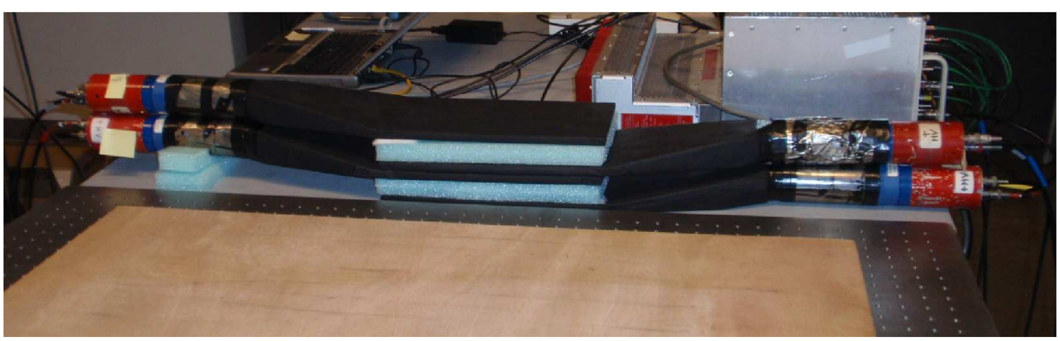

Figure 2: Stack of scintillator paddles used to count cosmic rays. Scintillator material is 12”×6”× 1cm. It is glued to a Lucite light pipe, which is wrapped in aluminum foil and heavy black paper. Each assembly is attached to a 10 stage PMT (Electron Tubes 9266KB). The high-voltage bias is between 800-1100 volts. High energy muons pass through all paddles, creating a set of coincident pulses.

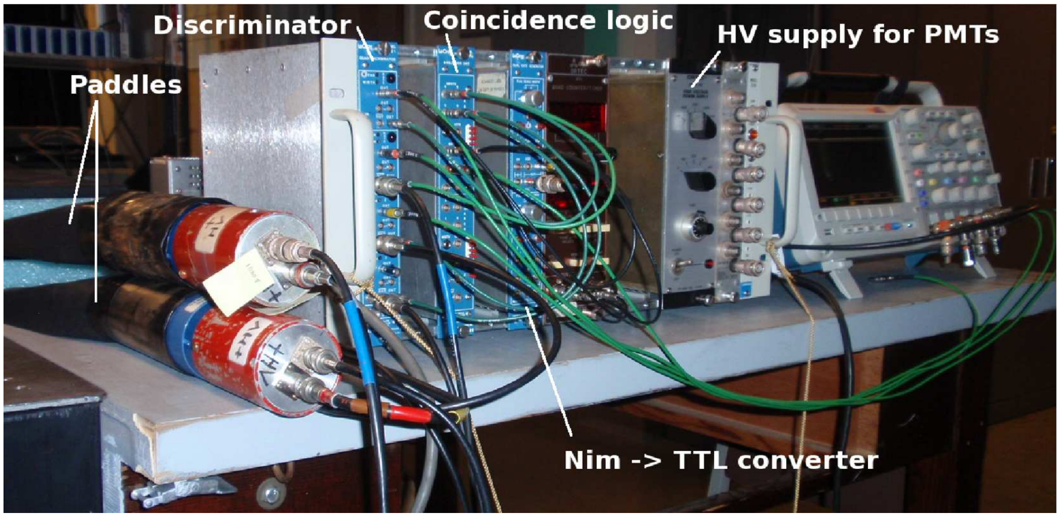

Figure 3: NIM electronics used to create pulses counted by data acquisition system. A 4-channel discriminator (LeCroy NIM 821) is set to pass pulses that are well above the noise level of background radiation. The discriminator feeds into a 4-fold logic unit (LeCroy NIM 365AL). When all 4 paddles fire simultaneoulsy, a pulse is created by the logic unit. The short negative “fast NIM” pulse is converted to a TTL-level “slow NIM” pulse by the gate generator (LeCroy NIM 222). The TTL pulse is used to gate one of the digital counters in the National Instruments DAQ card (NI USB-6221). The high volatge power supply (Ortec 456) supplies the bias to the PMT bases.

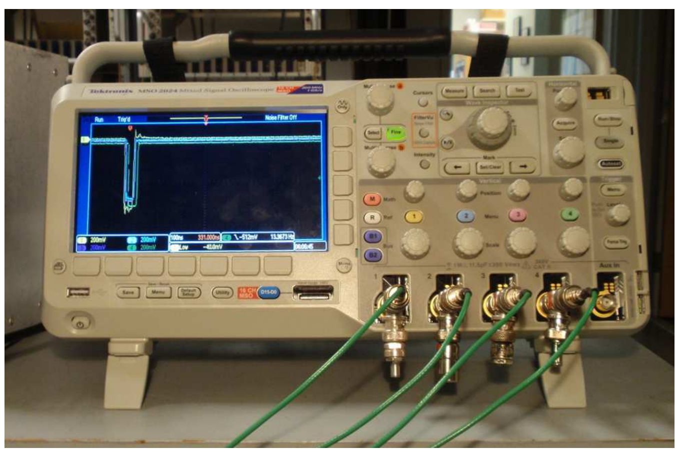

Figure 4: Oscilloscope showing sumultaneous fast NIM pulses, indicating the passage of a cosmic ray muon through all paddles.

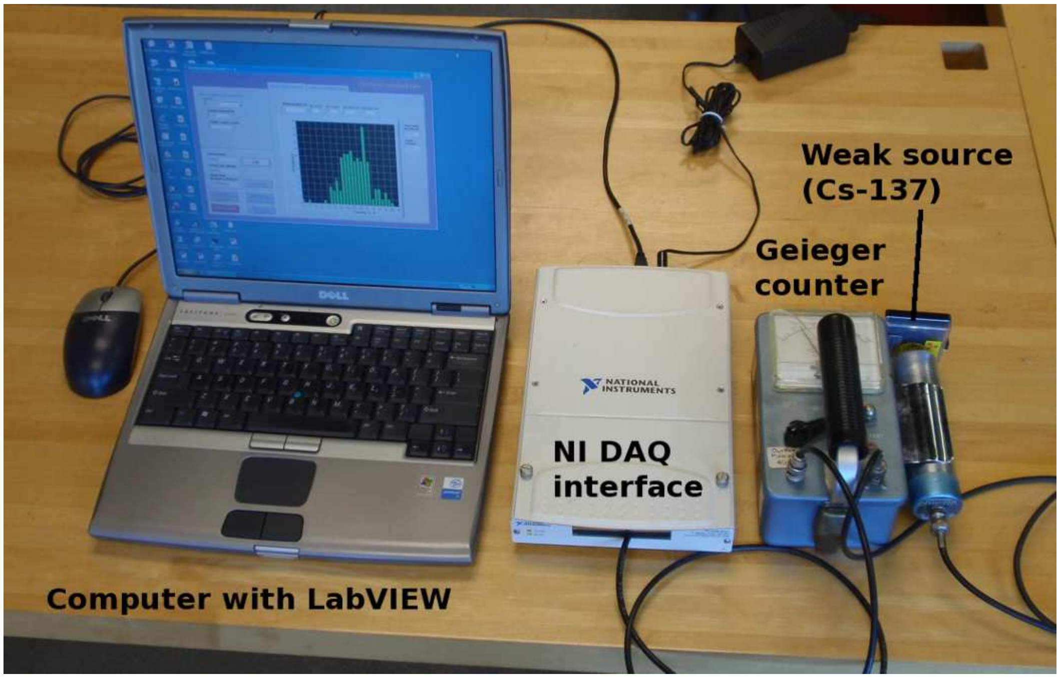

Figure 5: Minimal setup needed to count random pulses from a Geieger counter. The survey meter has a TTL level output that feeds directly into the counter gate on the DAQ interface. A computer with LabVIEW collects the data and the LabVIEW program allows analysis in a variety of ways. When used with the cosmic rays setup the output of the LeCroy 222 gate generator replaces the Geiger counter.

Equipment List

Apparatus for counting cosmic rays:

- Plastic scintillator paddles (3 or 4 preferred). Scintillator is Bicron BC-408, 1 cm slab cut into 12"×6" rectangle, which is glued to a Lucite light pipe (made in house). PMTs are Electron Tubes 9266KB, with a 2" circular face. PMT bases are made in house.
- Rack or holder for paddles to allow them to be stacked.
- NIM bin with additional 6 volt supply, e.g., Canberra model 2100.
- High voltage supply for PMTs, e.g., Ortec model 456. Useful to have more than one supply to account for variations in PMT sensitivity.
- NIM discriminator, e.g., LeCroy model 821 or Phillips model 704.
- NIM logic unit, e.g., LeCroy 365AL.
- NIM gate generator, e.g., LeCroy model 222. Needed to create TTL level pulse for Natioanl Instruments DAQ counter input.
- Cables to connect everything. High-voltage cables for PMT bias. 50-ohm cables for signals (RG-58 or RG 174). Note: signal cables from PMTs to discriminator inputs and from discriminator outputs to logic inputs must be the same length to insure that simultaneity is maintained for coincident events.
- Oscilloscope to study pule signals and set discriminator levels. Recommend 100MHz bandwidth or higher.
- Optional: NIM counter timer to do simple counting of pulse rates.

Apparatus for acquisition and analysis of pulse data:

- Computer with LabVIEW software. For the current experiment, computer is a Dell Latitude D600 running Windows XP with LabVIEW 7.1.1 and NI-DAQ version 8.8. The hardware uses the DAQmx library of functions.
- A multifunction DAQ interface, e.g., National Instruments USB-6221.

Apparatus to make random pulses, in lieu of cosmic ray counting setup:

- A Geiger counter with a TTL level output that can be used to generate pulses. The Pasco model SN-7927A could be adapted to this purpose.
- Weak radioactive source, e.g., Cs-137 available from PASCO Scientific.

Experiment Instructions and Tutorial

The attached sections provide copies of the instructions given to students in Physics 433 at the University of Washington:

1. *Cosmic Ray Counting.* Gives overveiw of cosmic rays, detailed information on setting up and using the paddle arrays, and a collection of exercises that explore the cosmic ray flux.
2. *Counting statistics of random events: A tutorial.* This tutorial presents derivations of some results from the theory of the statistics of random events that are studied in counting experiments.
3. Sample data.

Cosmic Ray Counting

Although we (almost) never think about it, we are bombarded by high-energy ionizing radiation every second of our lives. These high energy particles originate in outer space in the "solar wind". Very high in the atmosphere, the dominant particle species are protons or alpha particles, of extremely high energy (10^8 to 10^{20} eV). When these "primary" cosmic ray particles meet an atom in the upper atmosphere, the collision produces a "shower" of energetic hadrons, which then decay into energetic leptons. The most common of these "secondary" cosmic ray reactions are

$$\begin{aligned} \pi^+ &\rightarrow \mu^+ + \nu_\mu \\ \pi^- &\rightarrow \mu^- + \bar{\nu}_\mu \, , \end{aligned}$$

that is, the decay of pions into muons and neutrinos/antineutrinos.

A muon itself is unstable, and will decay (at rest) into an electron or positron (and associated neutrino or antineutrino) in a little over 2 microseconds. However, because of relativistic time dilation, the muons in high-speed flight from the upper atmosphere survive for much longer (according to our reference frame). Indeed, the lengthening of the muon lifetime in flight is a vivid example of relativistic kinematics in action. The discovery of the cosmic ray muon was one of the first indications that there could be particles other than the ones that make up ordinary matter: protons, neutrons, and electrons.

In this experiment, you will attempt to measure the flux of cosmic rays at about sea level (since Seattle is nearly there), and compare your value to a widely accepted one known since the 1940's. Cosmic rays come out of the sky in all directions, which means there are a couple of ways to talk about the flux. If we restrict the angle to straight down (zenith = 0), then the accepted intensity of high energy (hard) cosmic ray flux is

$$I_v = 0.82 \times 10^{-2} \text{cm}^{-2}\text{s}^{-1}\text{str}^{-1},$$

where the subscript v indicates that this is the vertically directed flux. This says that the number of cosmic ray particles passing through a vertically-facing surface is 0.82×10^{-2} particles per square centimeter per second per steradian. The steradian is a unit of solid angle in the way that a radian is a measure of planar angle. The solid angle intercepted by a sphere is 4π steradians, by the upper half of a hemisphere 2π steradians, and so forth. The steradian measure is necessary because the cosmic rays can come from all angles, not just straight down out of the sky.

If we integrate the intensity in the downward direction across a horizontal surface over the solid angle from horizon to horizon in all directions, we get the downward flux:

$$J = \int I(\theta) \cos\theta d\Omega = 1.27 \times 10^{-2} \text{cm}^{-2}\text{s}^{-1},$$

where the $\cos\theta$ term accounts for the fact that the flux across the surface from particles traveling at an angle to the surface normal will not be as large as the flux from particles traveling straight down.

If $I(\theta)$ were a constant, then the integral would give π times the intensity I_v, but it is notably less. This is because the flux intensity varies according to zenith angle. An approximate empirical relationship is

$$I(\theta) \approx I_v \cos^2\theta \, .$$

The muon telescope

Although cosmic rays come in a variety of particle types (muons, electrons, neutrinos, and their antiparticles) with a variety of energies, research has shown that most (about 80%) of the highest energy particles (the so called "hard" cosmic rays) to reach the surface of the earth are muons, both positive and negative. Thus, for the remainder of this write-up, I will use "muons" as a shorthand for "cosmic rays", since that is what we will mostly look at.

To measure the muon flux, we will use a "telescope", which in this context means an array of scintillator "paddle" detectors arranged with one scintillator overlapping another, as shown in Fig. 1. A muon traversing the array will, with some probability, cause a simultaneous pulse to occur in more than one detector, since with their high energy, they will usually not be stopped by a single paddle. Thus, we can use coincidence gating to count muon fly-bys as a way to distinguish these events from background radiation and photomultiplier noise.

Because the expected muon count rate is rather small, it is helpful to use a coincidence level of three or more in order to really suppress accidental coincidences. One of the aspects of this experiment you will explore is the rate of "accidentals" versus true coincidences. In addition, by comparing the rate of coincidences using 3 paddles to the rate using 4 paddles, you will be able to estimate the counting efficiency of the extra paddle, which you can then use to estimate a correction to your overall data set.

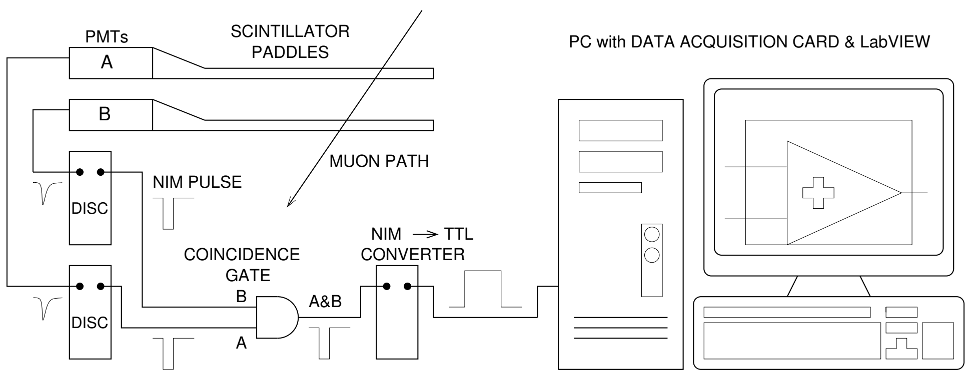

Figure 1: Schematic of the muon telescope apparatus. Although there are only two scintillator paddles shown in the figure, the actual setup uses between three and four paddles. The AND gate is a 4-fold NIM logic unit. Data are collected with the LabVIEW Interval Counter software.

Some details about the scintillator paddles: The scintillator material is Bicron BC-408 plastic scintillator cut into a sheet that is 1 cm thick, with an area of 12 inches by 6 inches. One end of the sheet is glued to a Lucite light pipe, which is then attached to the face of a 2 inch diameter PMT with optical grease. The PMTs are Electron Tubes model 9266KB. The paddle assembly is wrapped in aluminum foil, and then wrapped with heavy black paper and tape. A picture of an unwrapped paddle is shown in Fig. 2

Each paddle assembly rests on a shelf in a large wooden box. The shelves can be moved so that the vertical extent of the paddle array can be varied. The detectors are arranged so that the PMTs alternate which side of the paddle they are on—this minimizes the effect of a muon causing a coincidence by passing through multiple PMTs, rather than multiple scintillator paddles!

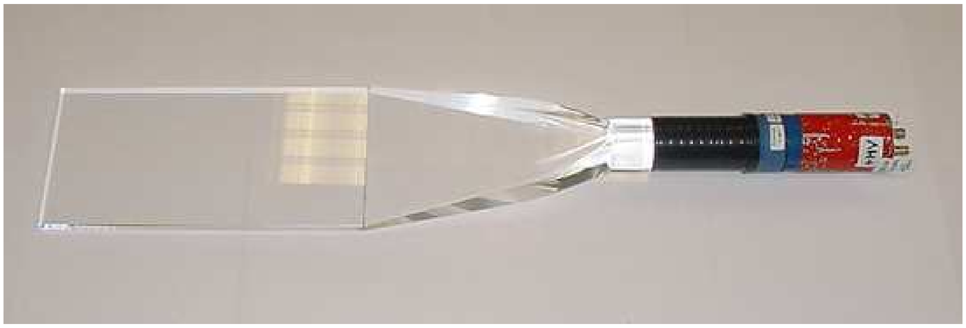

Figure 2: Scintillator paddle detector used in the muon telescope apparatus, before being wrapped with aluminum foil and black paper..

Setting up the telescope

The paddle array box will already be set up at your station with the paddles in place. You will need to connect the cables and set up the electronics.

First, note the tags attached to each PMT base indicating a voltage. This voltage has been chosen by the lab technician to produce a pulse of sufficient height when a ^{22}Na source is brought nearby. The variations in voltage from one paddle to the next are due to differences in the PMT manufacturing. (It is hard to make PMTs all exactly alike.)

Obtain SHV cables from the cable rack. (Use the long ones.) Since you have two high voltage supplies at your station, you must match, as well as you can, pairs of detectors to be plugged into each high voltage supply. For example, if your labels say 780V, 930V, 950V, and 960V, you would plug the 780 and 930 volt units into one supply, and the 950 and 960 units into the other.

Obtain 50ΩBNC cables from the cable rack. ***Important: the cables must all be the same length.*** (Remember, you are setting up coincidence counting.) Attach each cable appropriately, and double check to make sure the right HV cable goes to the right HV supply.

Start with one supply and paddle detector combination. Plug the signal cables into the two scope channels, appropriately terminated with 50Ω. Turn on the scope, turn the channel sensitivity all the way up (lowest V/div setting), and chose the appropriate trigger conditions. Slowly bring up the HV supply, following the usual protocol, while watching the scope. You are looking to make sure there are no light leaks or other problems. If you think you do see a problem, ask the instructor or TA for help.

Set the HV supply to the higher of the two settings recommended for your PMT pair. You will probably see a lot of noise pulses, and maybe some real event pulses too.

Obtain a ^{22}Na gamma source and put in in the holder. Carefully place the source somewhere between the two paddles that you are setting up, so that both are exposed to the source. Adjust the scope sensitivities so that you can comfortably see the pulses from both paddles. You may need to fiddle with the trigger settings somewhat. Estimate, as best you can from the scope screen, the peak height corresponding to the “bright line” of the 0.511 MeV annihilation photons for each detector in your pair. (No need to draw the pulse, just record its properties.)

When you are satisfied that your first pair of paddles is working OK, leave the HV on for them, and move on to the second pair of paddles. Repeat the procedure you just used: look at their outputs on the scope while bringing up the HV, check for problems, and check and record the response to the gamma source.

Feed the signal from each PMT into a discriminator channel. Use the source to set the width from each to 40 ns. After setting the width, set the discriminator level to the value you recorded for the “bright-line” peak height, or a bit less (i.e., 5–10 mV). Use the test point on the discriminator to dial in these numbers.

After setting the widths and the thresholds, rout the discriminator outputs into a channel in the 4-fold logic unit. Use a NIM counter/timer unit to measure 2-fold coincidence rates between the top two paddles, the middle two paddles and the bottom two paddles, over a reasonable counting period (10–20 seconds at most). You should see roughly the same count rates in each pair (with the appropriate pins set in the logic unit). If one pair seems excessively low, you my need to tweak the discriminator level and/or voltage to the PMT. Ask for help if you don't think it is working properly.

It is interesting to see what happens when you select 3-fold coincidence. You should see the count rate drop a lot, even though you still have the source nearby.

For your report. Explain why a 3-fold coincidence gives such a low count rate even with a ^{22}Na source in between two panels.

Characterizing the muon telescope

If all is well, you should be able to collect data in order to address two questions:

- What number of muons pass through the scintillators of the telescope? In other words, what is the *acceptance* of the telescope?
- What percentage of these muons does the telescope actually count? In other words, what is the *efficiency* of the telescope?

Acceptance of the telescope

The acceptance of the telescope depends on the surface area of the scintillator paddles, the solid angle subtended by surfaces of the paddle array, and on the angular dependence of the muon flux. Clearly a larger paddle area will see more muons, and it is not hard to believe that the counting rate would be directly proportional to this area, all else being equal.

The solid angle question is a bit more complicated. For a muon to make a coincidence between the top and bottom paddle, it must pass through both. Certainly a muon coming straight down will do this. But a muon coming in at an angle will also do this, if the angle is not too steep. The steepest angle would be a path from one edge of the top detector to the opposite edge of the bottom detector. If the top and bottom paddle are far apart, the solid angle will be smaller than if they are close together. However, calculating this quantity is tricky, for a couple of reasons. First, not every element of the detector sees the same solid angle. A small element of area dA in the center of the bottom paddle would register a count if the muon went through it and *any* spot on the top paddle, covering the whole surface A of that paddle. If you imagine a point on the bottom paddle, and rays intercepting this point and all points on the top paddle, you will get the solid angle for that portion of the array. But another element of the top paddle dA' located, say, near the edge, would see a different solid angle by the same construction. Moreover, the directions of the two solid angles, with respect to straight up (zenith angle $\theta = 0$) would also be different. It is known that the muon flux varies with zenith angle proportional to (about) $\cos^2\theta$. Thus, a good calculation of the acceptance of the telescope requires integrating the flux as a function of θ over the solid angle $\Omega(x, y)$ at each point (x, y) on the surface of the paddle, and then integrating this function over that surface A. At best, this is tedious; at worst, it requires a computer.

So, we will estimate the flux by making a couple of simplifying assumptions:

1. Assume that each small element of the paddle array sees the *same solid angle*, which is equal to the solid angle seen by the center of the paddle array.

2. Assume that the rectangular (approx 6"×12") paddles can be modeled as circular paddles having the same area.

The purpose of this second assumption is that it is easy to calculate what the flux per unit area would be for this cylindrical geometry. The solid angle would be that subtended by a cone. If the cone has zenith angle θ between the axis and the rim, then it is easy to show that

$$\int_{\text{cone}} d\Omega = \int_0^{2\pi}\int_0^{\theta} \sin\theta' d\theta' d\phi = 2\pi(1-\cos\theta)\ . \tag{1}$$

Likewise, to find the flux, we need to first integrate $I(\theta) \approx I_v \cos^2\theta$ over the solid angle. Hence,

$$J(\theta) \approx 2\pi \int_0^{\theta} \left(I_v \cos^2\theta\right) \cos\theta \sin\theta d\theta = \left(\frac{\pi}{2}\right) I_v(1-\cos^4\theta)\ . \tag{2}$$

As an exercise, you should verify that $J(\theta = \pi/2)$ gives the expected result of 0.0127 particles per cm^2 per second.

Use a ruler or tape measure to find the separation of the scintillator paddles. Also measure the area of the paddles, and check their alignment in the box by eye. You want to insure that they overlap as much as possible.

For your report. From your paddle dimensions, calculate the effective angle θ seen by cone whose area is the same as the area of one paddle and whose height is the same as the separation between the top and bottom paddle. Use this value, and the results above to calculate an estimate of the expected counting rate of hard muons. Comment on whether you think this estimate would be higher or lower than a more careful calculation. (Assume 100% efficiency of your detector paddles).

Efficiency of the paddles

Once you have your paddle array set up (HV on, discriminators set, test of coincidence between paddle pairs complete), remove the ^{22}Na source from the vicinity.

Connect the output of the logic to a gate generator or NIM/TTL converter, and feed the pulse into the computer data acquisition system. Run the LabVIEW "Interval Counter" program.

Assume that the paddles are labeled as follows, A for the top paddle, proceeding down to D for the bottom paddle. Use the NIM counter timer to perform the following tests:

1. Count with 2-fold coincidence using paddles A and D only.

2. Count with 3-fold coincidence using A, B, and D.
3. Count with 3-fold coincidence using A, C, and D.
4. Count with 4-fold coincidence using A, B, C, and D.

Make sure to collect enough counts so that you can estimate the counting rate with each configuration to better than 5% uncertainty, based on Poisson statistics. (How many counts is this?)

For your report. From your results of the above measurements, calculate the efficiency (and the uncertainty on that efficiency) for paddles B and C. Your efficiency estimates should be consistent with all of your measurements.

Describe the method and reasoning you used to find the efficiency of your paddles. Here is a definition of *efficiency*. The efficiency of a detector is the *detected* number of events divided by the *actual* number of events. In other words, if a detector counts 50 particles during a time when there were 100 particle hitting it that it *could* count, then the efficiency is 50% or 0.5. Here is something to think about: the efficiency is the probability of registering a count, given that there is a valid count to register. If we have two detectors arranged to count in a coincident manner, a particle may be detected by both detectors, by the first but not the second, by the second but not the first, or by neither. Only in the first case will a count be recorded, since coincidence is required.

Accidental coincidences

the efficiency analysis assumes that only real muons get detected by our array. But, as you will probably already have noticed, each paddle individually records a large number of counts per unit time. Only a small fraction of these are high-energy muons; most can be ascribed to PMT noise and low-level background sources. But with the low count rate for muons in this setup, the question of accidental coincidences becomes important.

As derived in the tutorial on counting statistics, the rate of accidental coincidences can be easily estimated, given the assumption that the coincidences are truly random. For two detectors, the rate of 2-fold coincidences is

$$R_2 = 2Tr_A r_B \ , \tag{3}$$

where T is the pulse width, and r_A and r_B are the rates of detectors A and B, respectively. For 3-fold coincidences the rate is

$$R_3 = 3T^2 r_A r_B r_C \ . \tag{4}$$

To test these relationships, we need to force any real coincidences out of the experiment. We can do this by adding in sufficient delay to the discriminator signals so that array is no longer in good synchronization.

Choose three paddles in your array (A, B, and C). Between the output of the discriminator and the input to the logic unit, add sufficient cable delay to move the second two paddles out of time with each other and with the first paddle. (for 40 ns pulse widths, 100 ns is probably sufficient.)

Use the pins on the logic unit to set up counting of the singles rates for each of the three detectors, and then count doubles rates (R_2) for any two detector pairs i.e., A+B, A+C, or B+C, and the finally count the triples rate R_3. You should see that the rate of random coincidences, especially for the triple, is much lower than when they are in time. This is a good way to show that you really are counting cosmic rays by the coincidence method!

For your report. Use your results to test the theoretical accidental rate predictions. Since these are statistical processes, your numbers will not match exactly, but you should be able to define a reasonable uncertainty based on Poisson statistics.

The long count

Reset correct timing in your paddle array in order to make a good coincidence count with 3 paddles. (You may use 4, if you find that their efficiency is very good.)

Start data collection with the Interval Counting program and let it run for a while. After about 5 minutes, stop data collection and look at the histograms for the following settings:

(a) Count distributions for interval times equal to 10 times the mean counts per minute.

(b) Count distributions for interval times equal to 20 times the mean counts per minute.

(c) Interval distributions for intervals of 1, 2, 5, 20, and 100 count intervals, using the "Interleaving" summing method on the interval distribution calculation.

After printing these histograms, restart the data collection, ***but do not clear the previous data.*** You want to see how the histograms build up over time. Repeat the above data collection after an additional 5 minutes has passed, and continue repeating the "5 minute interval/histogram grab" until the end of the lab period.

While the data are being collected, you can explore the settings on the program. What is the difference between "Scaling" and "Interleaving"? How can you explain the evolution of the histograms as you change the interval number? What is the effect of changing the bin number on the interval distribution?

Ideally, you should have a final data set corresponding to 20 minutes or more of counting data.

For your report. Discuss the relationships among the various histograms that you see. In particular, note how the mean and variance change, as the data sets get larger, and as you vary how you look at the same data set. Also how do the means and variances of the two different types of distributions (interval vs. count) compare with each other?

For your report. Calculate the flux measured by your telescope, and compare it to the ranges expected from your earlier calculation. Don't forget to calculate an uncertainty in the number. (It will probably be somewhat less.) From your efficiency measurements, make a correction to your data, by assuming that the efficiency of the outer two paddles is the same as the geometric mean of the efficiency of the inner two paddles. This will give you an estimate of the efficiency of the telescope as a whole. How much error do you get by ignoring the rate of accidental coincidences?

Optional: Simulation of random data

If you click on the "Show features" button, you will reveal a tab that will allow you to simulate data according to different interval distribution functions. Try this feature out. It is especially interesting to study the interval distribution evolution for the "uniform" function. Extra credit will be awarded according to how much you study and discuss.

Prepared by D. Pengra
counting_telescope.tex -- Updated 30 July 2008

Counting statistics of random events: A tutorial

This tutorial presents derivations of some results from the theory of the statistics of random events that are studied in counting experiments.

The statistics of counting random events falls under the rubric of what statisticians call "point processes". A point process in time is one where a particular type of event is classified only according to a single point on a time line. Some examples of phenomena that could generate time points are the births of babies in a particular hospital, the passage of cars past a counting gate on a toll bridge, or the arrival of particles from outer space at a detector. Such point processes may be further categorized as "Poisson processes" if the occurrence of any particular point is completely independent of the occurrence of any other point, as is the case in our cosmic ray experiment. If the criterion of independence is not satisfied, the point process may be more generally known as a "renewal process", a name which is derived from the study of part failures and replacements in industrial contexts (e.g., a new light bulb is not installed until the old one burns out) [4].

This section presents in a tutorial fashion the derivation of some important formulas in the theory of Poisson processes. These results are neither new nor little known, especially among statisticians. The goal is to show derivations that a student with the mathematical background of first-year calculus can follow and believe. These are "physicist's proofs" where a plausible argument is chosen over rigorous demonstration. Suggested exercises are given to help the student cement his or her understanding.

The raw data for our statistical investigations comes in the form of a random pulse train, as illustrated in Fig. 1. The pulse width T_w is important only insofar as it determines the maximum rate of pulses that may be represented by the pulse train, since pulses which occur more frequently than $1/T_w$ cannot be resolved. The occurrence time of each pulse, denoted as t_1, t_2, etc., as measured from some (arbitrary) starting time, is marked at the leading edge of each pulse.

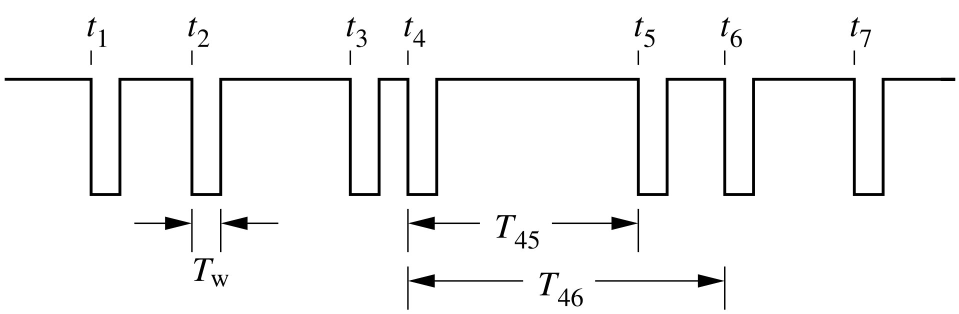

Figure 1: A random pulse train, as might be seen on an oscilloscope at a particular instant. This example shows negative-going ("fast NIM") pulses typical of those used in nuclear counting experiments.

If the pulse train is being created by the detection of decays from a long-lived radioactive nuclide or by cosmic rays, experiments show that although the time between specific pulses T_{ij} is completely unpredictable, the time intervals do converge to a well-defined average τ in the following sense: If one counts very many pulses (i.e., thousands) N over a long time interval T, then $\tau \approx T/N$. More specifically, if one estimates τ by counting N pulses over number of trials, the variation of the calculated τ's from these trials will cluster about each other with a fractional standard deviation of $1/\sqrt{N}$. (Later we will strengthen this assertion.) We can define an average *rate* of pulses r by the relation $r = 1/\tau$.

We are interested in two problems:

1. What is the *distribution of intervals* between successive pulses in a given pulse train? More generally, what is the distribution of *scaled intervals*: intervals between two pulses separated by an arbitrary number of pulses in between?
2. What is the *distribution of counts* within a succession of fixed-length periods?

To make these problems clear, we remind the reader of a few basic definitions. A *distribution* is the collected result of many *trials* of a particular experiment. Each trial produces a value of a *random variable*. For example, in the distribution of counts, one trial consists of adding up the number of counts that happen in a particular time period. This number is the value of the random variable $N(t)$, where t denotes the time length of the period. The collection of numbers $N(t_i)$, where i is an index denoting each trial period, is the distribution. Since $N(t)$ can only take on non-negative integer values, the count distribution is *discrete*. The random variable in the interval distribution is the time between pulses T_{ij}. Since it can take on any nonzero value, the interval distribution is *continuous*. Distributions are often presented as *histograms* (vertical bar graphs), in which the horizontal axis represents the value (or range of values, in the case of a continuous distribution) of the random variable and the height of the bars represents how often that value occurs.

The theoretical problem is the derivation of probability distribution functions: functions that give the likelihood of finding a particular value or range of values of the random variable. We will address the theoretical problem by making a basic assumption about our pulse train and invoking two rules. The rules are from the theory of probabilities [1]. Assume that two events of a similar type A and B follow from the same proximate cause. The probabilities of these events $P(A)$ and $P(B)$ are subject to these two rules:

> **I.** *The rule of compound probabilities.* If the two events A and B are mutually exclusive, then the probability of having either event $P(A \,\text{or}\, B)$ is given by the sum $P(A) + P(B)$.
>
> **II.** *The rule of independent probabilities.* If the two events A and B are independent of each other, then the probability of having both events occur $P(A \,\text{and}\, B)$ is given by the product $P(A) \times P(B)$.

Simple examples of these two rules are illustrated by considering rolls of dice. A die has six numbered sides, and when it is thrown any one number will be equally likely to land face up: the probability of any number is 1/6. If we ask, "What is the probability to roll an even number?", we are invoking rule **I**, since the roll of any number excludes the roll of any other number. Thus,

$$P(\text{even}) = P(2) + P(4) + P(6) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = \frac{1}{2} \, .$$

If we ask, "What is the probability to roll two 1s in a row?", we are invoking rule **II**, since each successive roll of the die is independent of any previous roll. Thus

$$P(1 \,\&\, 1) = P(1) \times P(1) = \frac{1}{6} \times \frac{1}{6} = \frac{1}{36} \, .$$

If we ask, "What is the probability of rolling only one 1 in two rolls of the die?", we want to know the probability of a 1 on the first roll *and* not 1 on the second *or* a 1 on the second roll *and* not on the first. Thus, since successive rolls of the dice are independent events and the two possible desired outcomes are mutually exclusive,

$$\begin{aligned} P(\text{only one 1 in 2 rolls}) &= P(1;1\text{st}) \times P^*(1;2\text{nd}) \;+\; P^*(1;1\text{st}) \times P(1;2\text{nd}) \\ &= \frac{1}{6}\times\frac{5}{6} \;+\; \frac{5}{6}\times\frac{1}{6} \\ &= \frac{10}{36}\,, \end{aligned}$$

where P^* is the compliment of P, that is, it is the probability of *not* obtaining the specified outcome. By convention, probability distributions are normalized: the random variable *must* take on one of its possible values, and since any value X either does or does not occur, $P(X) + P^*(X) = 1$, by rule **I**.

The basic assumption we make about our data is

> Any infinitesimally small interval of time dt is equally likely to contain a pulse. The probability of finding a pulse in dt is given by $r\,dt$.

For this assumption to make sense, $r\,dt$ must be much smaller than one, and the probability of finding two pulses in dt must be negligible. As long as r is a constant, these restrictions can be assured by making dt small enough.

We will first apply the rules we have stated to a real problem in particle physics: random coincidences.

Rate of random coincidences

The use of coincidence counting is pervasive in elementary particle experiments for one big reason: noise. For example, if you set up a scintillator paddle to count the passage of cosmic-ray muons, *most* of the counts you record will be due to non-muon events such as electrical noise, low-level background radiation, etc. Fortunately, muons have fairly high energy, so that if they pass through a thin scintillator paddle they will continue out the other side and can be detected by a second paddle nearby. Such an event produces two pulses, one in each detector, at (nearly) the same time. By using a logic gate, one can force the counting electronics to record only those pulses which occur in coincidence. A schematic of this setup is shown in Fig. 2. In order for the coincidence technique to work profitably, the rate of random coincidences must be appreciably lower than the rate of "true" coincidences. We can use our results to predict this rate.

Let the rate of random pulses recorded by detector A be be r_A and the rate of random pulses recorded by detector B be r_B. Each time a pulse from either A or B enters the gate, the gate is triggered: it "opens" for a short time interval T called the *resolving time*, meaning that if a pulse from the other detector occurs during T, then a coincidence pulse is produced at the output. Since r_A and r_B are the average rates of random pulses, the rate of coincidences will also be random (in the absence of any true coincidence events). Let R_2 be this rate; more specifically, we let $R_2\,dt$ be the probability that a random coincidence event between 2 detectors will occur within the infinitesimal time interval dt.

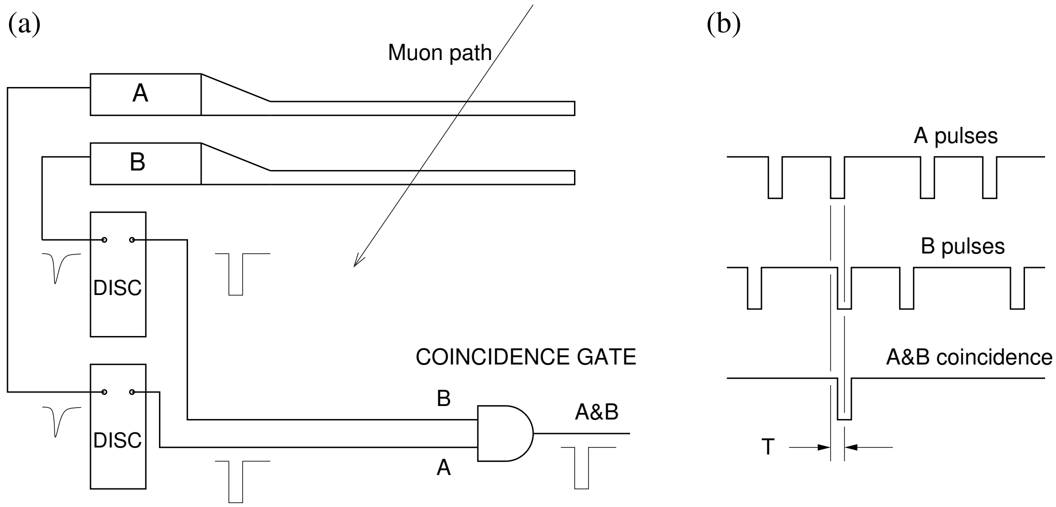

Figure 2: (a) A two-paddle coincidence setup. Variable height detector pulses are discriminated and turned into digital pulses. The coincidence gate produces a pulse when two input pulses overlap within a gate time T. (b) Example of two pulse trains showing one coincidence.

By rules **I** and **II** $R_2\,dt$ is given by the following construction:

$$R_2\,dt = \left(\begin{array}{l}\text{Probability of the gate}\\ \text{being triggered by}\\ \text{detector } A \text{ in } dt.\end{array}\right) \times \left(\begin{array}{l}\text{Probability of detector}\\ B \text{ delivering a pulse}\\ \text{within resolving time}\\ T \text{ of the trigger time.}\end{array}\right)$$
$$+ \left(\begin{array}{l}\text{Probability of the gate}\\ \text{being triggered by}\\ \text{detector } B \text{ in } dt.\end{array}\right) \times \left(\begin{array}{l}\text{Probability of detector}\\ A \text{ delivering a pulse}\\ \text{within resolving time}\\ T \text{ of the trigger time.}\end{array}\right) .$$

In symbolic form, this equation reads

$$R_2\,dt = r_A\,dt \times r_B T + r_B\,dt \times r_A T\ .$$

Typically the resolving time T is much shorter than the average rates r_A and r_B so that $rT \ll 1$. In this case, we may approximate the probability of finding a pulse in the finite time T by using the infinitesimal probability: $r\,dt \approx rT$. So the rate of 2-fold random coincidences R_2 is given by

$$R_2 = 2 r_A r_B T\ . \tag{1}$$

Exercise 1 *Show that if you added a third paddle C to the setup, the rate of random 3-fold coincidences R_3 would be equal to $3 r_A r_B r_C T^2$.*

Distribution of intervals

Now, let's return to our main questions. We first calculate the probability that there will be *no* pulse in a finite time interval T. Since we assume that a pulse may occur at any instant with equal

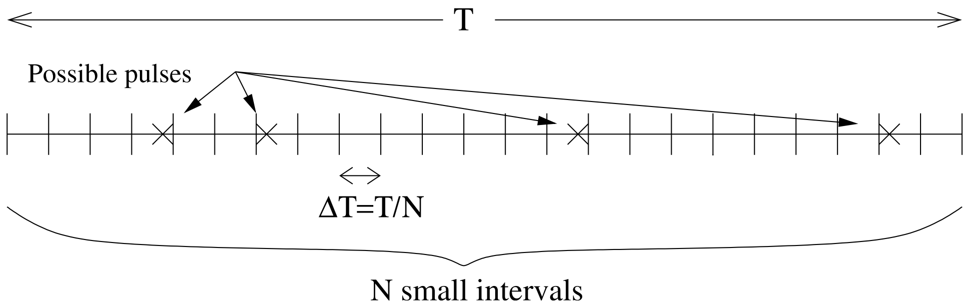

Figure 3: Construction used in calculating the probability of zero pulses in a finite time T.

probability, it does not matter where along the time line we start our clock. Since our assumption involves the infinitesimal interval dt, we divide T into small N small intervals ΔT, with an eye toward taking the limit of small ΔT. This construction is shown in Fig. 3. In any interval of length ΔT, when this interval becomes small, the probability of finding a pulse is $r\Delta T$, so by Rule **1**, the probability of not finding a pulse is $1 - r\Delta T$. Since the probability of finding a pulse in any small interval is independent of the probability of finding a pulse in any other interval, the probability of not finding a pulse in T, $P_0(T)$ is, by rule **II**, the product of the probabilities for all of the N small intervals:

$$P_0(T) \approx (1 - r\Delta T)^N \ . \tag{2}$$

If we write $\Delta T = T/N$, the limit of infinitesimal ΔT is found by letting N get large, and in this limit, the approximation tends to equality. Thus,

$$P_0(T) = \lim_{N\to\infty} \left(1 - \frac{rT}{N}\right)^N = e^{-rT} \ . \tag{3}$$

The final step in this unsurprising result can be proved by taking the logarithm of both sides of the equation, since $\ln(1 + x) \approx x$ for small x.

We now use the result of Eq. (3) to derive the probability density function of 1-pulse intervals, $I_1(t)$. We define this continuous probability density function as follows: $I_1(t)\,dt$ is the probability of finding a 1-pulse interval of length between t and $t + dt$. This means the probability of finding a pulse during the infinitesimal time interval dt after a time interval of length t following a previous pulse during which there have been no pulses. This situation is illustrated by Fig. 4.

If we pick a pulse in the train, and assign its time as $t = 0$, then we obtain $I_1(t)$ by the following construction from rule **II**:

$$\begin{aligned} I_1(t)\,dt &= \left(\begin{array}{l}\text{Probability of finding}\\ \text{zero pulses between 0}\\ \text{and } t\end{array}\right) \times \left(\begin{array}{l}\text{Probability of finding}\\ \text{a pulse between } t \text{ and}\\ t + dt\end{array}\right) \\ &= P_0(t) \times r\,dt \\ &= re^{-rt}\,dt \ . \end{aligned}$$

Thus

$$I_1(t) = re^{-rt} \ . \tag{4}$$

This result is known as the exponential distribution and has a couple of interesting features. First, the condition that each infinitesimal interval dt is equally likely to contain a pulse leads to the

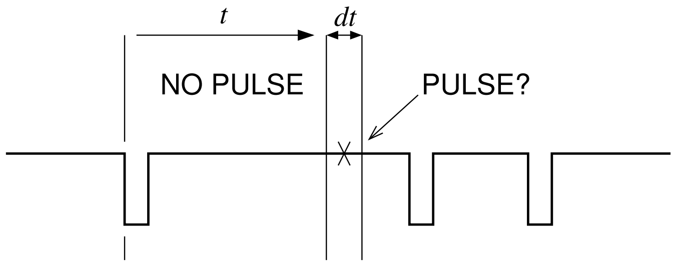

Figure 4: The construction leading to $I_1(t)\,dt$ asks how likely is a pulse to occur in dt after a previous pulse.

result that shorter intervals are much more common than longer ones. This is a vindication of the colloquial rule that "disasters come in threes": completely random events tend to cluster together in time. Perhaps more practical is the prediction that one may find the average rate r in a counting experiment by plotting the experimental interval distribution on a semi-log scale: both the slope and intercept will give r (after a some arithmetic).

Exercise 2 *(a) Show that $I_1(t)$ is normalized, that is, prove $\int_0^\infty I_1(t)\,dt = 1$. (b) Show that the average 1-pulse interval τ_1 is equal to $1/r$. The average interval is given by $\int_0^\infty tI_1(t)\,dt$.*

The situation for the distribution of 2-pulse intervals is a bit more complicated. We want the likelihood of finding an interval of a particular length t between two pulses that contains a single pulse anywhere inside at $t' < t$. One such possibility is shown in Fig. 5. The (differential) probability of finding an interval like this is given by rule **II**:

$$
\begin{aligned}
dP(t,t') = {}& \\
& \begin{pmatrix}\text{Probability of}\\ \text{finding 0}\\ \text{pulses between}\\ \text{0 and } t'\end{pmatrix} \times \begin{pmatrix}\text{Probability of}\\ \text{finding a pulse}\\ \text{between } t' \text{ and}\\ t'+dt'\end{pmatrix} \times \begin{pmatrix}\text{Probability of}\\ \text{finding 0}\\ \text{pulses between}\\ t' \text{ and } t\end{pmatrix} \times \begin{pmatrix}\text{Probability of}\\ \text{finding a pulse}\\ \text{between } t \text{ and}\\ t+dt\end{pmatrix} \\
= {}& P_0(t') \times r\,dt' \times P_0(t-t') \times r\,dt\ .
\end{aligned}
$$

This construction is nothing more than the probability of finding two 1-pulse intervals of particular lengths. But since the intermediate pulse can happen anywhere in the interval, by rule **I** we must sum the individual probabilities over all values of t' between 0 and t. Hence, we see that

$$
\begin{aligned}
I_2(t)\,dt &= \left(\int_0^t I_1(t')I_1(t-t')\,dt'\right) dt \\
&= \left(\int_0^t re^{-rt'}\,re^{-r(t-t')}\,dt'\right) dt \\
&= \left(r^2e^{rt}\int_0^t dt'\right) dt \\
&= r(rt)e^{-rt}\,dt\ .
\end{aligned}
$$

So the distribution function for 2-pulse intervals is

$$
I_2(t) = r(rt)e^{-rt}\ . \tag{5}
$$

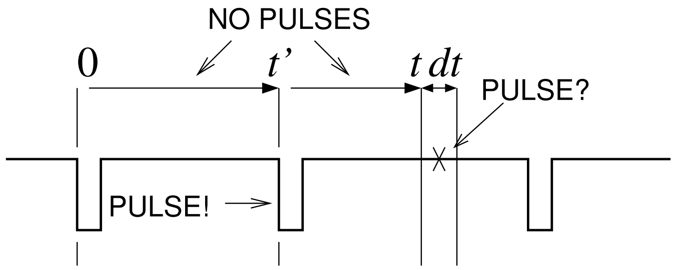

Figure 5: Construction used in calculating the probability a 2-pulse interval of length t.

Likewise, we may find the distribution for 3-pulse intervals $I_3(t)$ by multiplying $I_2(t')$ by $I_1(t-t')$ an integrating over all $t' < t$. This process leads to a general recursion relation

$$I_n(t) = \int_0^t I_{n-1}(t') I_1(t-t')\, dt' \ . \tag{6}$$

When applied to the exponential distribution the formula gives

$$I_n(t) = r\frac{(rt)^{(n-1)} e^{-rt}}{(n-1)!} \ . \tag{7}$$

Equation 7 is known as the (integer) gamma distribution or Erlang(ian) distribution (after the Danish engineer A. K. Erlang who used it to characterize telephone networks).

Exercise 3 *Convince yourself of the $(n-1)!$ factor in Eq. (7) by using $I_3(t)$ and $I_1(t)$ to show that*

$$I_4(t) = r\frac{(rt)^3 e^{-rt}}{1 \cdot 2 \cdot 3}$$

Exercise 4 *Show that the average n-pulse interval τ_n is equal to n/r, but that the most probable interval $T_{n,\mathrm{peak}}$ is given by $(n-1)/r$. The most probable interval is the interval at the peak of the distribution, and is found by evaluating $dI_n(t)/dt = 0$. Hint: $\int_0^\infty x^n e^{-x}\, dx = n!$.*

Graphs of $I_n(t)$ are shown in Fig. 6(a) for a few values of n. A couple of things are notable. First, the probability of finding zero length intervals for $n > 1$ is zero. This makes sense in that there must be at least one pulse within each such n-pulse interval, and even a single pulse takes time. More interestingly, the relative width of the peak of $I_n(t)$ shrinks as n increases. This is shown most clearly by scaling $I_n(t)$ so that the functions all have the same mean value $(1/r)$ while maintaining the normalized (unit) area, as is shown in Fig. 6(b). What this means is that for larger n, the collection of intervals becomes more uniform.

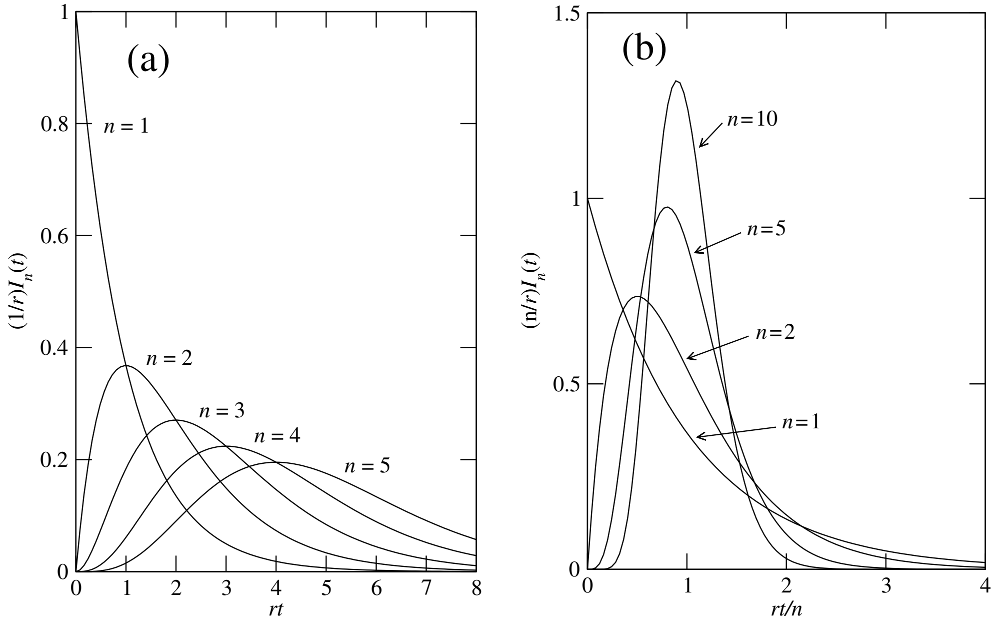

Figure 6: The n-pulse interval density functions, after Knoll [2, p. 98]. (a) $(1/r)I_n(t)$ for n from 1 to 5. Notice that the functions peak at $rt = n - 1$. (b) $(1/r)I_n(t)$ scaled by the transformations $t \to t/n$ and $(1/r)I_n(t) \to (n/r)I_n(t)$ to show that as n increases, the function becomes more sharply peaked around the mean value $\tau_n = n/r$.

Exercise 5 *Evaluate the variance σ_n^2 of $I_n(t)$. The variance is given by the formula*

$$\sigma_n^2 = \int_0^\infty (t^2 - \tau_n^2) I_n(t)\, dt \ .$$

Then show that the fractional standard deviation in τ_n, σ_n/τ_n is equal to $1/\sqrt{n}$.

Distribution of counts

Up to now, we have been framing our investigation thusly, "Given a fixed number of counts, how are the intervals that contain exactly this number distributed?" In other words, we are using the count number as a *parameter* in our problem, and letting the time be the *variable*. But frequently, our experimental situation is different: the independent parameter is the time period, and the dependent variable is the number of counts that occur within that time period. This is a related, but not identical question. Most obviously, as stated in the introduction, a time interval can take on any value, but a count number can only take integer values; the distribution of interval lengths is a *continuous* distribution, whereas the distribution of counts is a *discrete* distribution.

The relationship between the two distributions can be stated plainly: A period of a given length containing q counts, T_q, has *fewer* than n counts if and only if the interval containing n counts, T_n, is *longer* than T_q. That is,

$$q < n \text{ if and only if } T_n > T_q \ .$$

This suggests the following interpretation of what has been derived so far. The probability of finding an n-count interval of between length T and $T + dt$ is the same as the probability of finding

Table 1: A comparison of the two types of distributions that are contained in the "Poisson" formula. Note $\sigma_{_E}^2 = \tau_n^2/n$ and $\sigma_{_P}^2 = \overline{n}$.

	Erlang	Poisson
Type:	Continuous	Discrete
Variable:	t	n
Mean:	$\tau_n = n/r$	$\overline{n} = rt$
Variance:	$\sigma_{_E}^2 = n/r^2$	$\sigma_{_P}^2 = rt$

exactly $n-1$ counts in an interval of length T times the probability of finding one more count in T to $T + dt$. Stated in terms of our formulas:

$$I_n(T)\,dt = \frac{(rT)^{(n-1)}e^{-rT}}{(n-1)!} \times r\,dt = P(n-1; rT) \times r\,dt\ . \tag{8}$$

Thus, we identify $P(n-1; rT)$ as none other than the Poisson distribution giving the probability of finding $n-1$ events given that the average number of events is rT. This fact is used in the derivations of the interval distribution function given by Knoll [2] and Melissinos [3]. Most derivations of the Poisson distribution are based on taking a particular limit of the discrete binomial distribution. The derivation given here follows a more direct route from the basic rules of probability as applied to random pulse trains.

As is well known, the Poisson distribution is a distribution where the mean is equal to the variance; only one parameter is needed to define the whole distribution. This is not true of the Erlang distribution; for a given rate constant r and count number n, the mean interval time is n/r but the variance is n/r^2. A comparison of the two distributions given by our formula is given in Table 1.

References

[1] Wannier, Gregory H., *Statistical Physics,* (Dover, Mineola NY, 1987), p. 29.

[2] Knoll, Glenn F., *Radiation Detection and Measurement*, 2nd edition, (John Wiley, New York, 1989), pp. 96–99.

[3] Melissinos, Adrian C., and Jim Napolitano, *Experiments in Modern Physics*, 2nd edition (Academic Press, New York, 2003), pp. 401–404, 470–473.

[4] Cox, D. R., *Renewal Theory,* (Meuthen and Co. Science Paperbacks, Butler & Tanner Ltd., 1967).

Prepared D. Pengra

`counting_stats_tutorial_b.tex -- Updated 29 April 2008`

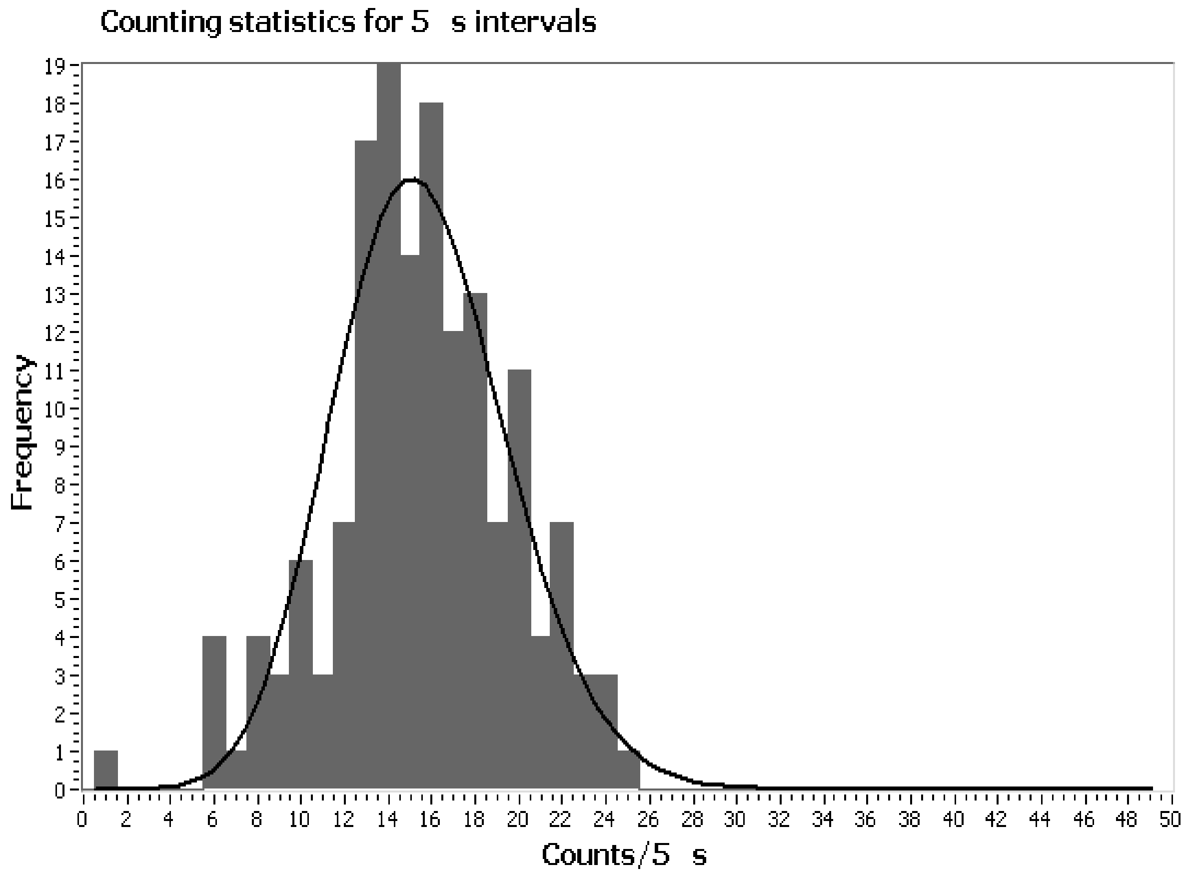

7/2/2013 4:16:09 PM
Total counts: 2448
Mean time between pulses: 0.32075 s
Mean counts/5 s: 15.586
Variance: 16.421
Histogram bins: 50
Curve: Poisson distribution

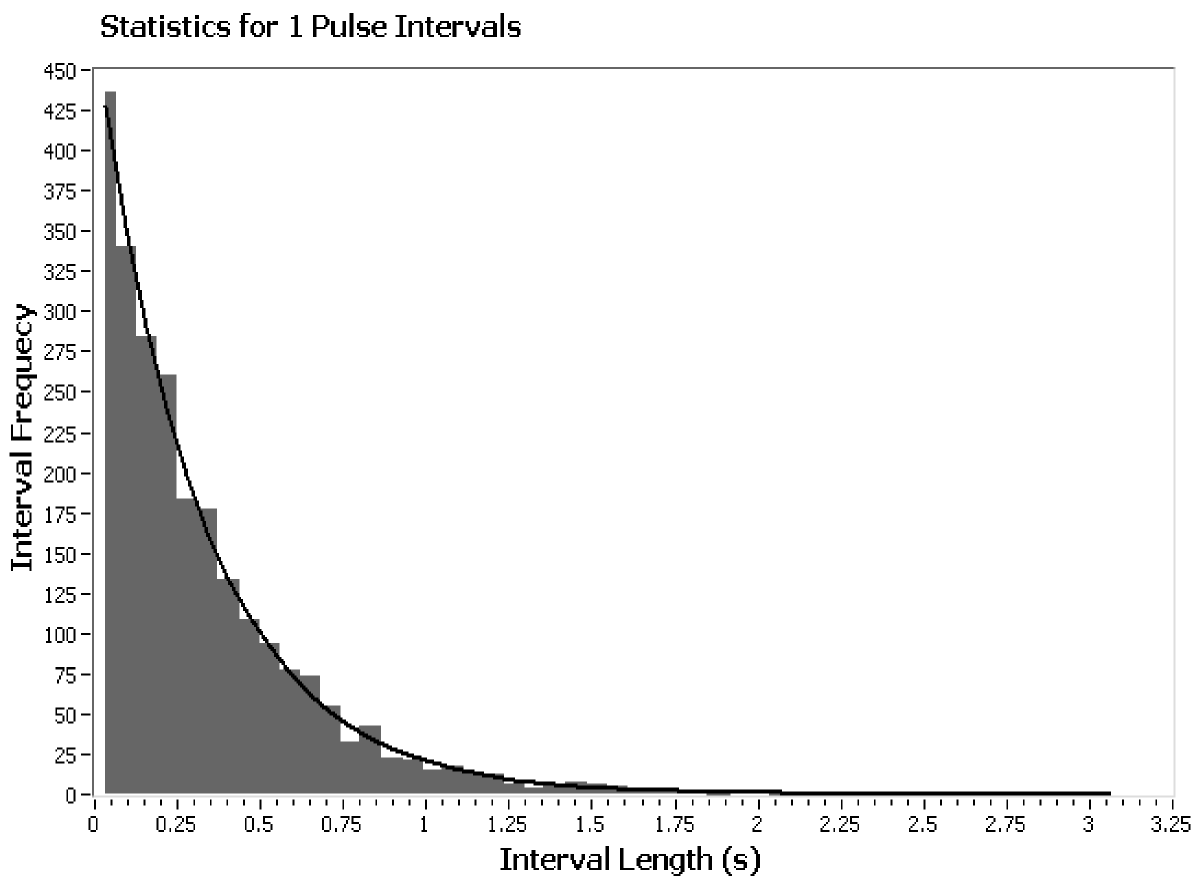

7/2/2013 4:21:35 PM
Total counts: 2448
Mean time between pulses: 0.32075 s
Scaled summing method
No. of histogram bins: 50
Mean interval for 1 counts: 0.321 s
Variance: 0.105 s^2
Curve: Exponential -> Erlang (empirical)

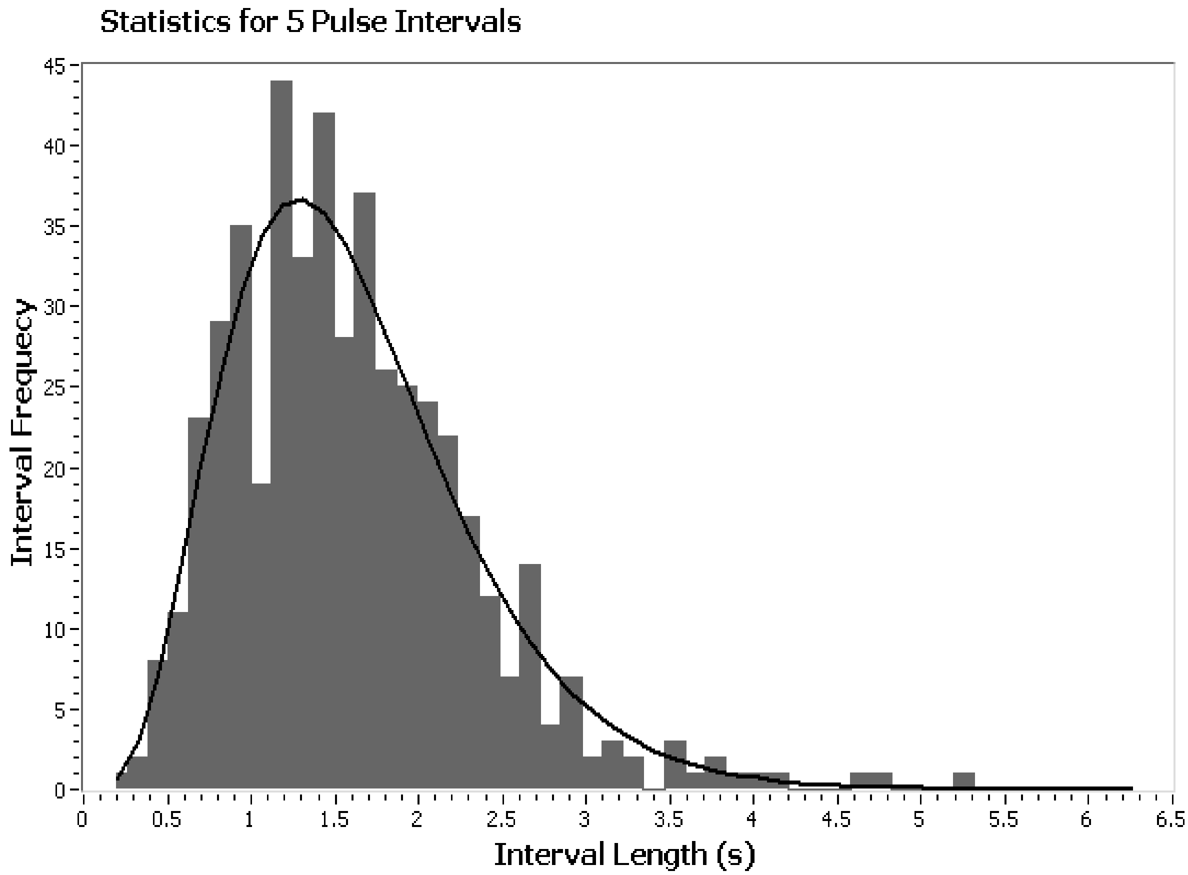

7/2/2013 4:22:43 PM
Total counts: 2448
Mean time between pulses: 0.32075 s
Scaled summing method
No. of histogram bins: 50
Mean interval for 5 counts: 1.602 s
Variance: 0.534 s^2
Curve: Exponential -> Erlang (empirical)